

How to Cope with Noise in the Iterated Prisoner's Dilemma

JIANZHONG WU

Chinese Academy of Sciences

ROBERT AXELROD

University of Michigan

Noise in the form of random errors in implementing a choice is a common problem in real-world interactions. Recent research has identified three approaches to coping with noise: adding generosity to a reciprocating strategy; adding contrition to a reciprocating strategy; and using an entirely different strategy, Pavlov, based on the idea of switching choice whenever the previous payoff was low. Tournament studies, ecological simulation, and theoretical analysis demonstrate (1) a generous version of tit-for-tat is a highly effective strategy when the players it meets have not adapted to noise; (2) if the other players *have* adapted to noise, a contrite version of tit-for-tat is even more effective at quickly restoring mutual cooperation without the risk of exploitation; and (3) Pavlov is not robust.

An important feature of interactions in the real world is that choices cannot be implemented without error. Because the other player does not necessarily know whether a given action is an error or a deliberate choice, a single error can lead to significant complications. For example, on September 1, 1983 a South Korean airliner mistakenly flew over the Soviet Union (Hersh 1986). It was shot down by the Soviets, killing all 269 people aboard. The Americans and Soviets echoed their anger at each other in a short, but sharp, escalation of cold war tensions (Goldstein 1991, 202).

The effects of error have been treated under the rubric of "noise." The best way to cope with noise has become a vital research question in game theory,

AUTHORS' NOTE: We thank John Harrington for valuable discussion. We thank the University of Michigan LSA College Enrichment Fund (R. A.) and the Chinese Fellowship for Scholarly Development of CSCC (J.Z.W.) for financial support.

JOURNAL OF CONFLICT RESOLUTION, Vol. 39 No. 1, March 1995 183-189
© 1995 Sage Publications, Inc.

especially in the context of the iterated prisoner's dilemma.¹ Clearly, when noise is introduced, some unintended defections will occur. This can undercut the effectiveness of simple reciprocating strategies. For example, Molander (1985) has shown that in the presence of any amount of noise, two tit-for-tat (TFT) players will, in the long run, average the same payoffs as two interacting RANDOM players.

Three different approaches to coping with noise have been proposed.

1. *Generosity*. Allowing some percentage of the other player's defections to go unpunished has been widely advocated as a good way to cope with noise (Molander 1985; May 1987; Axelrod and Dion 1988; Bendor, Kramer, and Stout 1991; Godfray 1992; Nowak and Sigmund 1992). For example, a generous version of TFT, called GTFT, cooperates 10% of the time that it would otherwise defect. This prevents a single error from echoing indefinitely.

2. *Contribution*. A reciprocating strategy such as TFT can be modified to avoid responding to the other player's defection after its own unintended defection. This allows a quick way to recover from error. It is based on the idea that one should not be provoked by the other player's response to one's own unintended defection (Sugden 1986, 110; Boyd 1989). The strategy called contrite TFT (CTFT) has three states: *contrite*, *content*, and *provoked*. It begins in content with cooperation and stays there unless there is a unilateral defection. If it was the victim while content, it becomes provoked and defects until a cooperation from other player causes it to become content. If it was the defector while content, it becomes contrite and cooperates. When contrite, it becomes content only after it has successfully cooperated.

3. *Win-Stay, Lose-Shift*. A completely different strategy can be used, one based on the principle that if the most recent payoff was high, the same choice would be repeated, but otherwise the choice would be changed. This strategy emerged from a simulated evolutionary process that included noise, but allowed strategies with memory of only the preceding move (Nowak and Sigmund 1993). Called Pavlov, it cooperates unless on the previous move it was a sucker (i.e., it cooperated but the other defected) or the other player was a sucker.

For completeness, we also analyze a fourth strategy, a generous version of Pavlov, called GPavlov. This strategy acts like Pavlov, but cooperates 10% of the time when it would otherwise defect.

1. Examples of recent theoretical and simulation studies of the noisy prisoner's dilemma and related games are Bendor, Kramer, and Stout (1991), Bendor (1993), Boyd (1989), Fudenberg and Maskin (1990), Godfray (1992), Kollock (1993), Lindgren (1991), Nowak and Sigmund (1992, 1993), and Young and Foster (1991). For a review of earlier work, see Axelrod and Dion (1988).

THE TOURNAMENT WITH NOISE

The basis for our analysis is the environment of the 63 rules of the Second Round of the Computer Tournament for the Prisoner's Dilemma (Axelrod 1984). These strategies provide a heterogeneous environment embodying a wide variety of ideas designed for doing well in the prisoner's dilemma game. The lengths of interactions vary, averaging 151 moves. To this environment we add 1% noise, meaning that for each intended choice there is a 1% chance that the opposite choice will actually be implemented. Although these rules were designed without regard to noise, they can still be used to provide useful setting for evaluating how new strategies will fare in a heterogeneous noisy environment.

The average score of each new rule when paired with the 63 rules of the tournament environment shows how well or poorly each does in a noisy environment.² The highest score is attained by GTFT which actually does better than any of the 63 rules submitted. CTFT also does very well, better than all but 5 of the 63 rules. Pavlov does poorly, ranking below 55 of the 63 rules. Adding generosity to Pavlov helps only a little: GPavlov ranks below 48 of the 63 rules.

To investigate the effects of different levels of noise, the 4 new rules were added to the 63 original rules, and the expanded tournament was run at various levels of noise from a tenth of a percent to 10%. Figure 1 shows the scores of the 4 new rules as a function of the noise level. The results show that at all levels of noise, GTFT and CTFT do well, and Pavlov and GPavlov do not. At the lower levels, GTFT is a little better than CTFT, but when noise is greater than 1%, CTFT is slightly better.

AN ECOLOGICAL SIMULATION

A more powerful test is to take into account that over time rules that are unsuccessful in the noisy environment are less likely than are relatively successful rules to be used again. A good way to do this is with an ecological analysis (Axelrod 1984, 48ff.). In an ecological analysis, the fraction of the population represented by a given rule in the next "generation" of the tournament will be proportional to that rule's tournament score in the previous generation. When this process is repeated over many generations, the

2. To assure stability of these results, the scores are averaged over 20 replications of the entire tournament.




Figure 1: Performance as a Function of Noise

proportion of the various rules changes, and the environment faced by each rule tends to emphasize those rules that have been doing relatively well in the noisy setting. The ecological simulation shows what happens when the rules that are ineffective in dealing with noise become a smaller part of the population, and those that are effective at dealing with noise become a larger proportion of the population.

The process begins with equal proportions of 67 rules: the 63 original rules and the 4 new ones. The noise level is set at 1%. The proportion of each rule is updated for 2,000 generations. Figure 2 shows the performance over time of the 6 rules that did best at the end of this process. R8, the rule that ranked eighth in the original tournament, did quite well over the first few hundred generations, but then slowly declined as the process continued to de-emphasize rules that were doing poorly in the noisy environment. By generation 1,000, CTFT was the leader. It continued to grow, eventually becoming 97% of the population at generation 2,000. GTFT had some early success, but then faded. Both versions of Pavlov declined to less than one part in a million as early as the hundred generation. The clear winner in this ecological simulation with noise was CTFT.

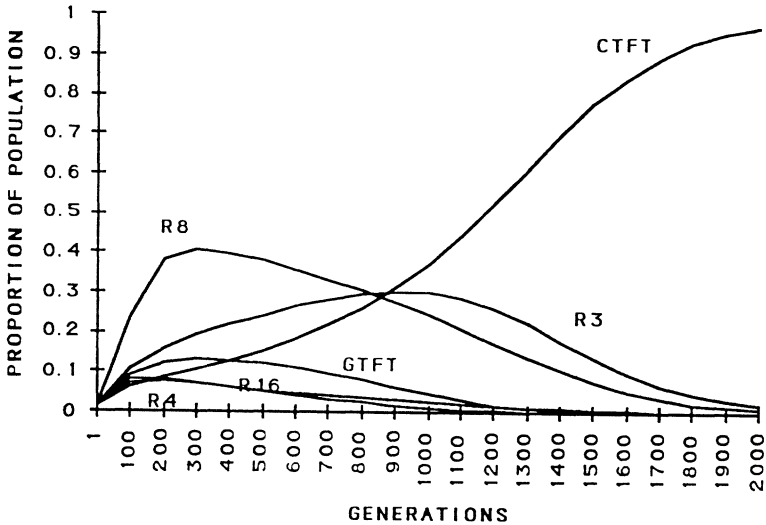


Figure 2: Ecological Simulation

NOTE: The strategies are the 63 original rules plus GTFT, CTFT, Pavlov, and GPavlov. R3 is the rule that ranked third in original tournament, R4 is the rule that ranked fourth, and so on. The results are shown for the top six rules in the 2,000th generation. The noise level is 1%.

STRATEGIC ANALYSIS

Both Pavlov and the contrite version of tit-for-tat have the desirable property that, when playing with their twin, they can quickly recover from an isolated error. If one of two Pavlov players defects because of an isolated error, both will defect on the next move, and then both will cooperate on the following move. If one of two CTFT players defects, the defecting player will contritely cooperate on the next move, and the other player will defect, and then both will be content to cooperate on the following move. Unfortunately for a player using Pavlov, its willingness to cooperate after a mutual defection can give the other player an incentive simply to defect all the time.³ The tournament and the ecological analysis both show that, although Pavlov may do well with its own twin, its success is not robust.

3. It pays always to defect when playing with Pavlov at low levels of noise if alternating T and P is better than always getting R. With standard notation, this is true when $T + wP > R + wR$ or $w < (T - R)/(R - P)$. Thus, with the standard payoffs of $T = 5$, $R = 3$, $P = 1$, $S = 0$, and any $w < 1$, it pays to always defect when playing with Pavlov.

Generosity is effective at stopping the continuing echo of a single error, whether the error was one's own or the other player's. The level of generosity determines how quickly an error can be corrected and cooperation restored. The problem is that generosity requires a tradeoff between the speed of error correction and the risk of exploitation (Axelrod and Dion 1988).

Contribution is effective at correcting one's own error, but not the error of the other player. For example, if CTFT is playing TFT, and the TFT player defected by accident, the echo will continue until another error occurs. Thus, in the original environment of 63 rules that were not designed to deal with noise, contribution was slightly less effective than generosity when noise was 1% or less. On the other hand, the ecological simulation showed that contribution is very effective as the environment becomes dominated by rules that are successful in the noisy environment. As the population becomes adapted to noise, contribution becomes more and more effective. In a population adapted to noise, correcting one's own errors is sufficient because the players one meets are also likely to be good at correcting their own errors.

CONCLUSION

In the presence of noise, reciprocity still works, provided it is accompanied by either generosity (some chance of cooperating when one would otherwise defect) or contribution (cooperating after the other player defects in response to one's own defection). Pavlov, a strategy based on changing one's own choice after a poor outcome, is not robust.

Generosity can correct an error by either player, but contribution can only correct one's own error. Thus, when the population of strategies one is likely to meet has not adapted to the presence of noise, a strategy like generous tit-for-tat is likely to be effective. On the other hand, if the strategies of the other players one is likely to meet have already adapted to noise, then a strategy like contrite tit-for-tat is likely to be even more effective because it can correct its own errors and restore mutual cooperation almost immediately.

REFERENCES

- Axelrod, R. 1984. *The evolution of cooperation*. New York: Basic Books.
 Axelrod, R., and D. Dion. 1988. The further evolution of cooperation. *Science* 242:1385-90.
 Bendor, J. 1993. Uncertainty and the evolution of cooperation. *Journal of Conflict Resolution* 37:709-34.

- Bendor, J., R. M. Kramer, and S. Stout. 1991. When in doubt: Cooperation in a noisy prisoner's dilemma. *Journal of Conflict Resolution* 35:691-719.
- Boyd, R. 1989. Mistakes allow evolutionary stability in the repeated prisoner's dilemma. *Journal of Theoretical Biology* 136:47-56.
- Fudenberg, D., and E. Maskin. 1990. Evolution and cooperation in noisy repeated games. *American Economic Review* 80:274-9.
- Godfray, H.C.J. 1992. The evolution of forgiveness. *Nature* 355:206-7.
- Goldstein, J. 1991. Reciprocity in superpower relations: An empirical analysis. *International Studies Quarterly* 35:195-209.
- Hersh, S. M. 1986. *The target is destroyed: What really happened to Flight 007 and what America knew about it*. New York: Random House.
- Kollock, P. 1993. An eye for an eye leaves everyone blind: Cooperation and accounting systems. *American Sociological Review* 58:768-86.
- Lindgren, K. 1991. Evolutionary phenomena in simple dynamics. In *Artificial life II: Proceedings of the workshop on artificial life*, edited by C. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen. Redwood City, CA: Addison-Wesley.
- May, R. M. 1987. More evolution of cooperation. *Nature* 327:15-7.
- Molander, P. 1985. The optimal level of generosity in a selfish, uncertain environment. *Journal of Conflict Resolution* 29:611-8.
- Nowak, M., and K. Sigmund. 1992. Tit for tat in heterogeneous populations. *Nature* 355:250-3.
- . 1993. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the prisoner's dilemma game. *Nature* 364:56-8.
- Sugden, R. 1986. *The economics of rights, co-operation and welfare*. Oxford: Basil Blackwell.
- Young, H. P., and D. Foster. 1991. Cooperation in the short and in the long run. *Games and Economic Behavior* 3:145-56.